

a' Referring back to Figure 3a, after initialization, alternative routes leading from the alternative-route-enabled nodes are computed (308). The alternative routes are computed along the initial route computed in step 330 of Fig. 3b to the destination device 110. Information associated with the alternative route is stored locally at the alternative-route-enabled nodes so an alternative route is readily available for rapid resumption of network traffic when a failure occurs in the initial route. Packets are forwarded using the initial routes (310). According to the stored topology information, each node forwards packets to the IP address of the next node on the specified link. In a MPLS framework, each node swaps a label on the IP packet before sending it to the next hop. While in session, RSVP hello extensions are implemented at each node providing hello packets that can be sent directly to other nodes on the initial route. A failure can thus be quickly detected (by a failure to respond to a hello message). Once a failure is detected (315), and if the detecting node is alternative-route-enabled (322), traffic for the failed route is directed to the pre-computed alternative route stored locally (320). Thereafter, packets are forwarded using the alternative route. If the detecting node is not alternative-route-enabled, a message is sent to an ingress node (324), the first node in the system on the initial route. Thereafter, the process ends (326) and the ingress node may invoke a back-up mechanism. In another implementation, when the node detecting a failure is not alternative-route-enabled, a failure message can be forwarded upstream to a nearest in-line alternative-route-enabled node on the initial route. The alternative-route enabled node can implement the alternative route and continue forwarding packets.

---